EXHIBIT 12

**IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
MARSHALL DIVISION**

TQ DELTA, LLC, **Plaintiff,** v.	§ § § § §	**JURY TRIAL DEMANDED**
COMMSCOPE HOLDING COMPANY, INC., COMMSCOPE INC., ARRIS INTERNATIONAL LIMITED, ARRIS GLOBAL LTD., ARRIS US HOLDINGS, INC., ARRIS SOLUTIONS, INC., ARRIS TECHNOLOGY, INC., and ARRIS ENTERPRISES, LLC,	§ § § § § § § §	**Civil Action 2:21-cv-310-JRG (Lead Case)**
NOKIA CORP., NOKIA SOLUTIONS AND NETWORKS OY, and NOKIA OF AMERICA CORP. **Defendants.**	§ § § § §	**Civil Action No. 2:21-cv-309-JRG (Member Case)**

DECLARATION OF DR. TODOR COOKLEV IN SUPPORT OF PLANTIFF'S OPENING CLAIM CONSTRUCTION BRIEF

I, Todor Cooklev, hereby declare under penalty of perjury:

1. My name is Todor Cooklev, and I have been retained by complainant TQ Delta LLC (collectively, "TQ Delta").

2. I am over eighteen years of age, of sound mind, and am not otherwise disqualified from making this Declaration. I have personal knowledge of the facts contained in this Declaration. I am competent to testify on the matters set forth below if I am asked to do so.

3. I have been asked to prepare this declaration in connection with the above captioned District Court actions, and have been asked to investigate and opine on issues relating to certain claim limitations of certain patents within the Patents-in-Suit, as discussed below, in addition to providing a relevant technical background.[1] I have been asked to explain what a person of ordinary skill in the art would have understood certain terms in the Patents-in-Suit to mean in and around the earliest effective filing dates. An explanation of my opinions is set forth below.

4. I base this Declaration on information currently available to me. I reserve the right to continue my investigation. I further reserve the right to expand, modify, and/or supplement this Declaration and my opinions as additional information becomes available to me, including in response to matters raised by the defendants and/or defendants' expert(s), or in view of relevant orders and findings by the Court.

[1] I understand that the Patents-in-Suit include the following U.S. Patent Nos. 7,453,881 ("the '881 Patent"); 7,570,686 ("the '686 Patent"); 7,844,882 ("the '882 Patent"); 8,090,008 ("the '008 Patent"); 8,276,048 ("the '048 Patent"); 8,462,835 ("the '835 Patent"); 8,468,411 ("the '411 Patent"); 8,495,473 ("the '5473 Patent"); 8,594,162 ("the '162 Patent"); 8,595,577 ("the '577 Patent"); 8,937,988 ("the '988 Patent"); 9,094,348 ("the '348 Patent"); 9,014,193 ("the '193 Patent"); 9,154,354 ("the '354 Patent"); 9,300,601 ("the '601 Patent"); 9,485,055 ("the '055 Patent"); 9,547,608 ("the '608 Patent"); 9,894,014 ("the '014 Patent"); 10,044,473 ("the '4473 Patent"); 10,409,510 ("the '510 Patent"); 10,567,112 ("the '112 Patent"); 10,833,809 ("the '809 Patent").

5. In this declaration, I cite to various documents and testimony. These citations are meant to be exemplary, and not exhaustive. Citations to documents or testimony are not intended to signify that my conclusions or opinions are limited to the cited sources or supported by the cited sources only.

6. I am being compensated as an independent consultant in this matter at the rate of $400 per hour for analysis of documents and preparation of any declaration or report. This compensation is not dependent on my opinions or testimony, or the outcome of this litigation.

I. QUALIFICATIONS AND BACKGROUND

7. My qualifications are stated more fully in my curriculum vitae, a copy of which is attached hereto as Exhibit 1.

8. I am currently Professor of Electrical and Computer Engineering at Purdue University in Fort Wayne, Indiana, where I have had several faculty and administrative positions. I have received research funding from the National Science Foundation (NSF), the Defense Advanced Research Projects Agency (DARPA), the U.S. Air Force Research Laboratory, the Office of Naval Research, several major companies, and other smaller technology companies. I have given and continue to give seminars, tutorials, and presentations worldwide. I was inducted into the Purdue Inventors Hall of Fame (2019).

9. I received a Ph.D. in Electrical Engineering from the Tokyo Institute of Technology in 1995. I have authored and co-authored more than 100 peer-reviewed articles. I am a named inventor on thirty-two (32) U.S. patents, most of which relate to the multicarrier communication systems. (A list of my publications and patents appears in my *curriculum vitae* attached as Exhibit 2014.) Further, I teach classes that cover DSL technology including, for example, the ECE 428 Communication System course which is among the rotation of courses

that I teach each year.

10. Some of my publications relate to the mathematics of Reed-Solomon (and more generally BCH) codes. For example, in the 1990s I co-authored a paper in the IEEE Transactions on Circuits and Systems entitled “Generalized Fermat-Mersenne number-theoretic transforms”, among other publications relating to Galois fields.

11. In addition to my academic experience, I have experience in technology and product development in the computer networking and data communications industry. My work has been in digital signal processing, software, and integrated circuit design for communication systems. Between 1999 and 2002 I worked on and contributed to the implementation of the DSL standards that existed at the time on an integrated circuit.

12. During that time, I was also working on improving these DSL standards. I was a participant in the T1E1.4, International Telecommunications Union (ITU) SG 15 Q4 committees. I have attended meetings of these committees and I have prepared, submitted, and presented documents relating to various technical matters considered by these committees.

13. Further, I have served in leadership roles. For example, I chaired the ad-hoc session on coding for DSL at the ITU-T SG15 Q4 meeting in Antwerp, Belgium in June 2000. The ad-hoc session produced a document BA-108R1 that defined the requirements towards the different coding proposals before the committee.

14. I have also designed and implemented low-density parity check (LDPC) encoders and decoders for DSL systems. LDPC is another type of coding used in communication systems. I presented my work on LDPC for ADSL at two standards committee meetings of ITU-T SG15 Q4, held in Antwerp, Belgium and Goa, India, in 2000.

15. I have also participated in other standards committees in the field of

communication systems. My long record of service and leadership includes serving as Chair of the IEEE Standards in Education Committee, a committee jointly sponsored by the IEEE Standards Association and the IEEE Educational Activities Board. Since January 2017 I have served as on the Editorial Board of the IEEE Communications Standards Magazine, a scholarly journal specifically devoted to the field of communication standards.

16. In 2020 I was elected to serve on the Board of Governors of the IEEE Standards Association (IEEE SA) for one term beginning January 2021. The Board of Governors provides overall leadership of the IEEE Standards Association. As a Board member I participate in the committee overseeing the products and services provided by the IEEE SA. I also contribute to the relationship between IEEE SA and other standards bodies worldwide.

17. I have served as an expert on patent litigation cases involving various communications technologies. A list of litigations in which I have served as an expert appears in my curriculum vitae.

18. I am qualified by education and experience to testify as an expert with respect to subject matter in the fields of protocols and standards for communication systems, hardware and software implementations, and interoperability among devices.

II. INFORMATION AND MATERIALS CONSIDERED

19. In forming my opinions expressed in this declaration, I have considered and relied upon my education, background, and experience. I have also reviewed the relevant Patents-in-Suit and their file histories as necessary to express the opinions in this declaration. I have also reviewed the materials cited in this declaration.

III. LEVEL OF ORDINARY SKILL IN THE ART

20. A person of ordinary skill in the art at or before the time of the inventions of the

Patents-in-Suit addressed in this declaration would have had a bachelor's degree in electrical engineering with 2-3 years of experience in DSL communication systems. Alternatively, the person would have had a master's degree in electrical engineering. I am familiar with the knowledge and capabilities of one or ordinary skill in this area based on my experience working with industry, with undergraduate and graduate students, with colleagues from academia, and with engineers practicing in the industry.

IV. APPLICABLE LEGAL STANDARD

21. For the purposes of this Declaration, counsel has instructed me to make the following assumptions:

a. Claim construction is solely a matter for the Court to decide and, in general, the ordinary meaning of the claim terms used in the patent to one of ordinary skill in the art is determined in the context of the patent's specification and the file history.

b. A "person of ordinary skill in the art" is a hypothetical person who is presumed to have known the relevant art at the time of the invention. Factors that may be considered in determining the level of ordinary skill in the art may include: (1) type of problems encountered in the art; (2) prior art solutions to those problems; (3) how quickly innovations are made; (4) sophistication of the technology; and (5) educational level of active workers in the field. In a given case, every factor may not be present, and one or more factors may predominate. A person of ordinary skill in the art is a person of ordinary creativity. A person of ordinary skill in the art would have the capability of understanding the scientific principles applicable to the pertinent art.

c. Claims are construed from the perspective of a person of ordinary skill as of the effective filing date of the patent application. For the '346 patent, that date is May 7, 2001.

d. Persons of ordinary skill in the art are deemed to read the claims in the context of the entire patent, including the specification and prosecution history. In other words, the terms are not considered in a vacuum. In the context of claim construction, the specification has been called "the single best guide" to the meaning of the claim terms.

e. Claim terms should be given their ordinary and customary meaning within the context of the patent in which the terms are used, *i.e.*, the meaning that the term would have a person of ordinary skill in the art in question at the time of the invention in light of what the patent teaches.

f. The plain and ordinary meaning is determined from the language of the claims, the specification, and the prosecution history of the patent at issue.

g. In construing a claim term, one looks primarily to the intrinsic patent evidence, which includes the patent abstract, specification, claims, and figures, and its prosecution history.

h. Extrinsic evidence may also be useful in interpreting patent claims when the intrinsic evidence itself is insufficient.

i. The usual and customary meaning of a claim term can be altered by the patent applicant if they choose to act as their own "lexicographer" and clearly set forth in the patent a different meaning of a claim term.

j. The meaning of a claim term can also be altered during the patent examination process by a clear and unequivocal disclaims and disavowals by the patent applicant makes about the meaning or scope of the term, and that such statements are recorded in the prosecution history of the patent application.

k. If a claim term is ambiguous or unclear, the term must be construed to determine how a person of ordinary skill in the art would have resolved in light of the rest of the patent specification, patent claims, and the application's prosecution history.

l. A claim is not indefinite, as long as it, having been read in light of the intrinsic evidence, informs one of skill in the art at the time of the invention about the scope of the invention with reasonable certainty.

m. It is improper to import limitations from embodiments in the specification.

n. It is also improper to import limitations from other parts of the claims thus rendering a claim term duplicative.

22. It is also improper to import additional or different language from other independent claims that would render such claims superfluous.

V. BACKGROUND OF THE TECHNOLOGY

23. This section describes the background technology for the Patents-in-Suit that are referred to as the "Family 2,"[2] "Family 3,"[3] "Family 9,"[4] and "Family 10"[5] Patents.

24. The following sections provide an overview of the field of the inventions and the subject matter of the Family 2, 3, 9, and 10 Patents. To facilitate an understanding of the inventions, I describe the state of the art and concepts important to understand the inventions.

A. Overview of Communications Systems

[2] I understand that the Family 2 Patents include the '881, '193, '601, and '014 Patents.

[3] I understand that the Family 3 Patents include the '882, '048, '5743, '608, and '510 Patents.

[4] I understand that the Family 9 Patents include the '411, '577, '348, '055, '4473, and '809 Patents.

[5] I understand that the Family 10 Patents include the '354 and '988 Patents.

25. In the process of data communication between two communication endpoints, a first (near end) and second (far end) transceiver for example, communicate one or more bits of data (data bits) are communicated serially. The two communication endpoints are commonly known as a customer premise equipment (CPE) transceiver and a central office (CO) transceiver.

26. The term "transceiver" is a combination (called a portmanteau in English grammar) of two words: transmitter and receiver. A transceiver can perform both the transmitting and receiving operations.

27. Typically, the transmitter portion of the transceiver and the receiver portion of the transceiver share at least some common circuitry (*e.g.*, memory or processor). For example, the Family 3 Patents illustrate that "transceiver 100 includes a transmitter portion 200 and a receiver portion 300," along with a shared processing module (110) and shared memory (120). *See* '890 Patent at 4:41-43.

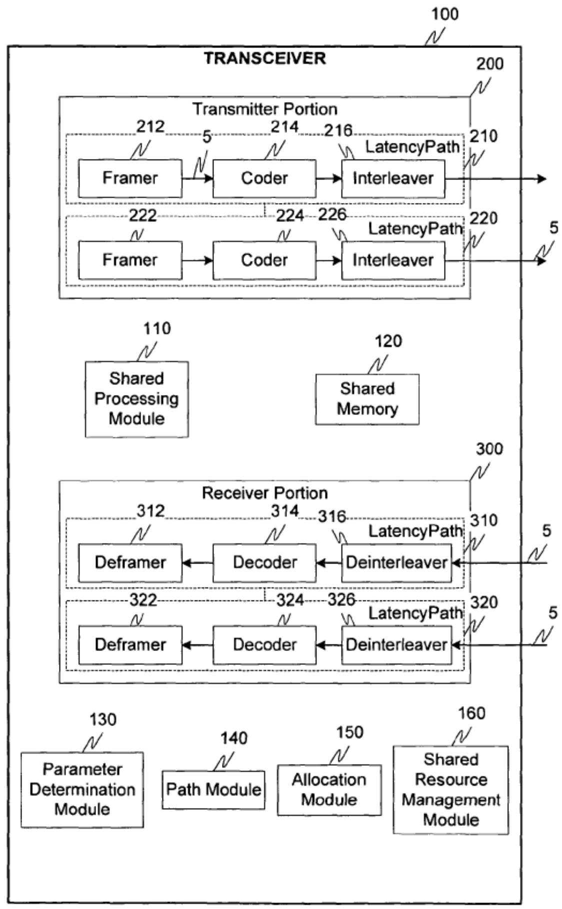

Fig. 1

'890 Patent at Fig. 1.

28. The data bits are communicated over a communication medium that may be, for example, a twisted-pair transmission line (also known as a phone line). The stream of data bits is communicated between the communication endpoints using symbols transmitted serially over a communication medium, such as a twisted-pair transmission line. Each symbol represents one or more bits. In DMT (which is an acronym for discrete multitone) systems, one DMT symbol can represent hundreds or thousands of data bits. An illustration of a bit stream is shown below. Each white box represents a single bit or a group of data bits.

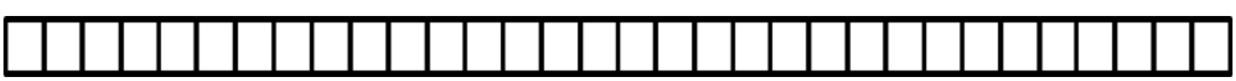

29. Digital communications systems typical transmit data as packets, which generally

refers to any basic data unit, such as a grouping of bytes. See ’956 Patent at 10:38-45 (“The term ‘packet’ includes any basic data unit, i.e., a grouping of bytes. For example, a packet could be an IP packet, an Ethernet packet, an ATM cell, a PTM packet, an ADSL Mux-Data frame, a PTM-TC codeword, an RS Codeword, a DMT symbol, or, in general, any grouping of data bytes or information.”).

30. Data communication between two communication endpoints is accomplished by transmitting bits of data. The two communication endpoints are commonly known as a customer premise equipment (CPE) transceiver and a central office (CO) transceiver. The data bits are communicated between the communication endpoints using symbols transmitted serially over a communication medium, such as a twisted-pair transmission line. Each symbol represents one or more bits. In DMT systems, one DMT symbol can represent hundreds or thousands of data bits.

31. Typically, the data includes payload data and overhead data. The payload data represents the bulk of the information for transmission between the CO and CPE, and is composed of, inter alia, web-page data, video, voice, and other data types. Overhead data is used by the transceivers to exchange various messages required by the protocol, including status and control data, etc.

32. Multicarrier transmission systems provide high speed data links between communication endpoints over twisted pairs (i.e., a telephone line). The twisted pair communication channel can also be used for other services. Indeed, multicarrier modulation can be used to communicate data over existing twisted pair communication channels that were not originally intended for data communications and may be subject to substantial electromagnetic interference (i.e., “noise”) and other impairments. For example, DSL systems must operate alongside existing Plain Old Telephone Service (POTS) on a twisted-pair transmission lines in

diverse and difficult conditions.

33. The Family 2, 3, 9, and 10 patents relate to communicating data using digital subscriber line (DSL) communications systems. Three main examples of these systems are Asymmetric DSL (ADSL), Very-High-Bit-Rate DSL (VDSL), and Fast Access to Subscriber Terminals (G.fast).

34. ADSL devices are also known as ADSL transceivers or ADSL Transceiver Units (ATUs). There are two main types of ATUs: ATU-R at the remote or customer side, and ATU-C at the Central Office side. An ADSL connection consists of an ATU-C connected to an ATU-R using a phone line. In the case of ADSL, the ATU-C and the ATU-R comprise the communication endpoints. In the case of VDSL, the functionally equivalent units are known as VTU-O at the central office side and VTU-R at the remote or customer side. In G.fast, the functionally equivalent units are known as FTU-O at the central office side and FTU-R at the remote or customer side.

35. Any communications system consists of three entities: transmitting entity, receiving entity, and channel entity. The transmitting entity and the receiving entity are on different sides of the channel entity. In the case of ADSL, VDSL, and G.fast these sides are also known as “CPE side” (Customer Premise Equipment side) and “CO side” (Central Office side). Equipment that is capable of transmitting and receiving (and therefore consisting of a transmitting entity and a receiving entity) are referred to as “transceivers” as described above.

36. DSL systems typically use multicarrier modulation. Multicarrier transmission systems provide high speed data links between communication endpoints over twisted pairs (i.e., a telephone line). The twisted pair communication channel can also be used for other services. Indeed, multicarrier modulation can be used to communicate data over existing twisted pair

communication channels that were not originally intended for data communications and may be subject to noise and other impairments. For example, DSL systems can operate alongside existing Plain Old Telephone Service (POTS) on a single twisted-pair transmission line without disruption.

37. Communications channels have finite capacity and the data rates cannot be arbitrarily high. As with any communication medium and communication protocol, there are limits to the data carrying capacity (i.e., data rate) of a DSL connection over a single phone line.

B. Addressing Noise, Interference, and Attenuation in Communications Systems

38. In practical communications systems not all bits of data that are transmitted will be received correctly. Bits of data communicated over the communication medium may be corrupted by noise or interference. Noise is defined as "unwanted disturbances superimposed upon a useful signal that tend to obscure its information content." IEEE 100 The Authoritative Dictionary of IEEE Standards Terms, 7th ed. (2000) at p. 732. Thus, bits of data transmitted over a communication medium in the presence of noise may not be correctly received.

39. There are several types of noise, including thermal noise, noise from adjacent telephone lines, and noise from wireless or radio transmissions. These types of noise tend to cause random bit errors. Another type of noise is impulse noise. Impulse noise is defined as noise that is "characterized by transient disturbances separated in time by quiescent intervals." IEEE 100 The Authoritative Dictionary of IEEE Standards Terms, 7th ed. (2000) at p. 732. One source of impulse noise is the switching of high-voltage electrical equipment. Impulse noise is intermittent and causes bursts of errors, i.e. consecutive errors.

40. In DSL-based systems, the noise and interreference is very significant because the communications medium (*e.g.*, telephone lines) are very noisy channels, especially with regard to impulse noise. It is therefore critical, as a technical matter, to efficiently detect and correct

errors that result from channel noise and interference.

1. Error Detection Techniques

41. Communications systems include a technique, typically based on Cyclic Redundancy Check (CRC), to determine if an error has taken place. Specifically, to permit detection of errors caused by impulse noise or interference, error detection bits, cyclic redundancy checksum ("checksum" or "CRC") bits for example, may be appended to groups of data bits before transmission. The transmitting end point may compute the checksum bits for the group of data bits and append it to the group of data bits prior to transmission. The figure below illustrates an exemplary bit stream that includes checksum bits 'C' for groups of data bits. Different shades of grey are used to differentiate the different groups of bits.

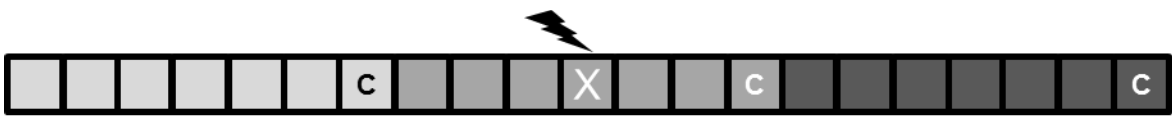

42. The receiving endpoint uses the checksum bits to confirm that respective groups of data bits were correctly received. For example, the receiving endpoint may independently compute the checksum for a group of data bits and compare the computed checksum with the checksum appended by the transmitting endpoint to the group of data bits. If the computed checksum matches with the appended checksum, the group of data bits was correctly received. Otherwise at least one of the data bits of the group was not correctly received, for example, because it was corrupted by impulse noise. In the example illustrated above, at least one data bit of the second group of bits is corrupted by impulse noise (marked with an 'X'). In this case the checksum computed for the second group of data bits will not match with the appended checksum computed by the transmitting endpoint.

2. Retransmission

43. One way to handle errors is to retransmit the data that has been corrupted. Retransmissions are not generally desirable because re-sending the same data more than once reduces the effective data throughput.

44. In retransmission schemes disclosed in the Family 9 Patents, after a first transceiver transmits a packet to a second transceiver, the transmitting transceiver stores the packet in its memory, i.e., retransmission memory or buffer, until the receiving transceiver communicates to the transmitting transceiver that the packet was correctly received or was not received/received with errors, or until the transmitting modem infers that the packet was not correctly receive or infers that it was correctly received. See ’956 Patent at 12:24-30 (“The low-PER packets can be stored for a sufficient amount of time to wait for a retransmission message from the receiver . . . . During this time, the transmitting transceiver can continue to receive packets from one or more higher layers, label these packets, if needed, and store these packets, ...”); *id.* at 12:33-40 (“[f]or successful retransmission, the receiving modem should be able to inform the transmitting modem which packet, or packets, need to be retransmitted.”); *See also* ’956 Patent at 14:38-62 (“For example, the receiving modem . . . can send an acknowledgment (ACK) message to the transmitting modem for every correctly received message or every predetermined number of packets.”); *id.* at 14:63-15:20 (“[a]lternatively, or in addition, a timeout value could be specified for the transmitting modem. This timeout value could correspond to the amount of time that the transmitting modem should wait for an ACK for particular packet before retransmitting the packet. The timeout value could be set to be at least as long as the round-trip delay required for the transmitting modem to send a packet to the receiving modem and for the receiving modem to send an ACK back to the transmitting modem. If an ACK is not received by the timeout value, the transmitting modem could retransmit the packet. Alternatively, or in

addition, a negative acknowledgment (NAK) could be sent to the transmitting modem when a packet is detected as errored or missing.”).

45. To reduce packet errors, if the transmitting modem understands that a packet was not correctly received, the transmitting transceiver retransmits the packet that is stored in the retransmission buffer. *See* ’956 Patent at 2:10-12 (“the low-PER packets can be stored in a retransmission buffer, e.g., memory, that can be used to reduce packet error.”).

46. The Family 9 Patents, in describing the amount of memory required to implement a retransmission scheme, explains:

> The transmitting modem must store a packet for retransmission for a time equal to the round-trip delay from when the packet is sent to when the retransmission message is received. During this time the transmitting modem continues to receive packets from the higher layer and continues to store these packets in the same way. Therefore, the storage requirements in octets can be computed as:
>
> $$MinimumTXmemory(octets) = roundtripdelay * datarate,$$
>
> where the roundtrip delay is the time equal to the round trip delay from when the packet is sent to when the retransmission message is received, and the data rate is the data rate of the connection that is transferring the packets.

’956 Patent at 16:52-67.

47. The Family 9 Patents further explain that to implement the disclosed retransmission functionality, the receiving modem includes comparable functionality to the transmitting modem. A14 ('956 Patent) at 13:20-25 ("The receiving modem, e.g., receiver PTM-TC, which in this case is illustrated as the transceiver 300 and includes comparable functionality to that described in relation to transceiver 200."). Notwithstanding the fact that the receiving transceiver receives retransmitted packets, the Family 9 Patents refer to the functionality implemented in the receiving transceiver that facilitates retransmission also as retransmission

functionality (for example, "re-receive" is not used to refer to the functionality in the receiving transceiver).

48. For the receiving transceiver, the minimum retransmission memory requirement is similar to that of the transmitting transceiver:

> For the receiver, the minimum receiver storage requirements can be determined in a similar manner. More specifically, the RX PTM -TC must store a packet before passing it to the higher layer for a time equal to the round trip delay from when a retransmission message is transmitted to when the retransmitted packet is received. This is equal to storage requirements in octets (same as transmitter):
>
> $$MinimumRXmemory(octets) = roundtripdelay * datarate,$$
>
> where the roundtrip delay is the time equal to the round trip from when a retransmission message is transmitted to when the retransmitted packet is received and the data rate is the data rate of the connection that is transferring the packets.

’956 Patent at 17:10-23.

49. More specifically, the RX PTM -TC must store a packet before passing it to the higher layer for a time equal to the round-trip delay from when a retransmission message is transmitted to when the retransmitted packet is received. This is equal to storage requirements in octets (same as transmitter).

50. To reduce the negative impact of noise, error correction schemes may be used. The techniques that correct errors without the need for retransmissions are called Forward Error Correction (FEC). Use of Forward Error Correction causes latency (i.e., delay), but typically less delay than retransmission. The Family 9 Patent specification explains that “packets containing, for example, video information (such as IPTV) may have the requirement for a very low packet error rate.” ’956 Patent at 1:65-66. To provide a low packet error rate or to negate the effect of impulse noise, error correction schemes may be used. Exemplary error correction

schemes include retransmission and interleaving/forward error correction (FEC). '956 Patent at 17:63-18:5.

3. Forward Error Correction and Interleaving

51. Forward error correction (FEC) and interleaving is another scheme that provides protection against impulse noise. In an exemplary forward error correction scheme, data bits are encoded into corresponding sets of codewords prior to transmission, where the codeword comprises data bits plus a number of parity or redundancy bits. The encoding method used is one that allows the data bits to be recovered from its corresponding codeword. The codewords are generated using a suitable encoding scheme, Reed-Solomon (RS) coding for example.

52. Use of Reed-Solomon coding is one Forward Error Correction technique used with DSL systems. Reed-Solomon coding is a type of block coding. In Reed-Solomon coding, groups of data bits may be encoded into corresponding sets of codewords prior to transmission. In place of the data bits themselves, the groups of data bits in each codeword are then transmitted sequentially over the communication medium. The groups of bits prior to encoding can be referred to as "symbols" or "data symbols" (*e.g.*, a symbol of 8 bits, which are different from the multicarrier symbols that are being transmitted). A group of bits obtained after encoding is a "codeword" that is encoded per the technique (and then decoded upon receipt to recover the original symbols).

53. The encoding method used is one that allows a group of data bits to be recovered from its corresponding codeword. The figure below illustrates the groupings of data bits and the corresponding codeword generated from each group. Groups of bits and codewords would typically be communicated sequentially, for example as illustrated below, group 1a from codeword 1 being communicated first in time and group 4d from codeword 4 being communicated last in time:

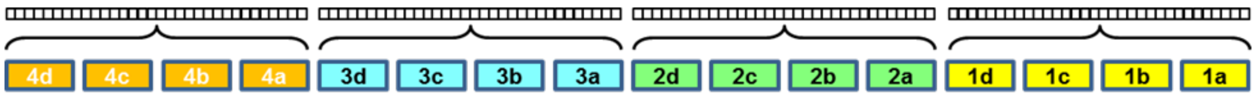

54. Reed-Solomon (RS) codes are codes that operate over symbols of s bits each (*e.g.*, a symbol of 8 bits would result in s being 8). The major parameters of RS codes are codeword length n and message length k (both typically measured in the number of symbols). Coding involves two processes: encoding and decoding. Encoding is performed at the transmitter and decoding is performed at the receiver.

55. The technique adds "parity" symbols to each message to form a codeword. The parity symbols add redundancy to the transmission to allow the receiver to still decode the transmission even if there are errors (up to a maximum amount).

56. On the encoding side, the encoder takes k data symbols and then fills up the remainder of the codeword length (n) with parity symbols (this result in the number of parity symbols being n-k). The image below shows the full codeword (of length N) with the data symbols and the parity symbols. This code is denoted as RS (n, k).

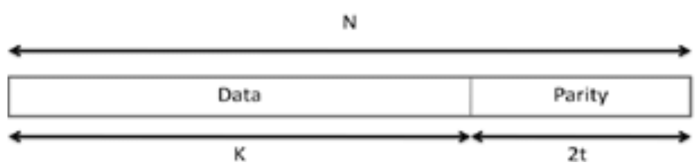

57. The maximum number of corrupted symbols that a decoder can correct in a given codeword is referred to as t, which is half the number of parity symbols included in the codeword (which mathematically is $\frac{n\text{-}k}{2}$).

58. Thus, for any coding technique, including Reed-Solomon, the size of a codeword is greater than the size of the data that the codeword encodes.

59. For example, the ITU-T recommendation for ADSL G.992.1 specifies that the

number of parity bytes per RS codeword is 0, 2, 4, 6, 8, 10, 12, 14, or 16. For example, one RS code is RS (255, 239). This code operates using 8-bit or 1-byte symbols. Each codeword contains 255 bytes and consists of 239 message bytes and 16 parity bytes. The decoder can automatically (i.e. without retransmission) correct up to $(n-k)/2=8$ bytes of errors. These 8 erroneously received bytes can be located anywhere among the 255 bytes of the codeword.

60. Even using a forward error correction scheme (such as Reed-Solomon coding), the transmission may still be sufficiently corrupted through impulse noise that the original data bits cannot be recovered. This is because impulse noise is intermittent and causes bursts of errors, i.e. consecutive errors, that can exceed the maximum threshold that the forward error correction scheme allows (e.g., in the example above, more than 8 erroneously received bytes among the 255 bytes of the codeword). As example of impulse noise is shown below, with the corrupted bits shaded in red and showing that portions of codeword 3, including all bits of symbol 3c and some of the bits of symbols 3b and 3d, are corrupted by impulse noise.[6]

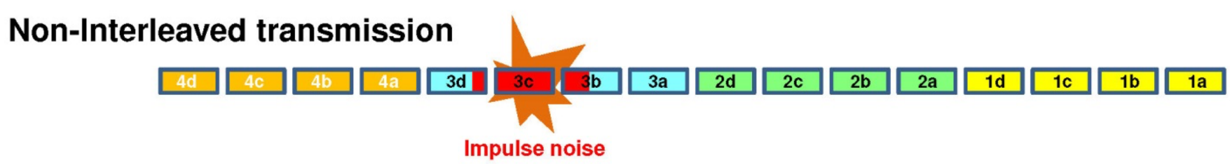

61. How much an impulse noise event impacts the ability to receive the underlying data depends on how much error the forward error correction scheme (such as Reed Solomon encoding) can correct for and if the noise exceeds that maximum value.

62. If one or more bits from a codeword are corrupted by noise, the receiving endpoint may nevertheless be able to recover all of the original data bits using the portions of the

[6] The illustration below refers to a "non-interleaved" transmission. I describe the interleaving process below as a technique to reduce the impact of impulse noise by using an "interleaved" transmission.

codeword that are correctly received. To allow recovery of the original data bits from a codeword, a minimum or threshold number (or percentage) of bits from the codeword must be correctly received. If, however, more than a threshold number of bits from the codeword are corrupted by noise, the original data bits represented by a codeword will not be recoverable. This minimum or threshold number is a function of the level of encoding performed when translating the group of data bits to the codeword.

63. Because impulse noise involves short, intermittent bursts, a number of consecutive bits spanning a number of symbols/bytes of a codeword may be corrupted while other adjacent bits are received correctly (corrupted bits are shaded in red in the illustration above). If the duration of the impulse noise is sufficiently long, then more than a threshold number of bits of a codeword (for example, more than one quarter or 25% of the bits in a codeword) may be corrupted, which would make it impossible to recover the data bits corresponding to that codeword. In the example illustrated above, portions of (blue) codeword 3 (all of bits in byte/symbol 3c and some of the bits of byte/symbols 3b and 3d totaling about 40% of the bits of the blue codeword) are corrupted by impulse noise. In this simple example where the threshold is 25%, it would be impossible to recover the original data bits that were used to generate (blue) codeword 3.

64. An improvement to the above-described error correction scheme spreads the impact of impulse noise across multiple codewords by using interleaving. ’956 Patent) at 17:47-48 (“Interleaving with RS coding is an effective way of correcting channel errors due to, for example, impulse noise.”); A722 (ITU-T Rec. G.993.2 (02/2006) at p. 51) (“Interleaving shall be provided in all supported latency paths to protect the data against bursts of errors by spreading the errors over a number of Reed-Solomon codewords.”). This approach involves manipulating

multiple codewords by rearranging the order in which bytes of the multiple codewords are transmitted. Bytes of a codeword, that otherwise would have been contiguous (i.e., transmitted sequentially), are rearranged so that once contiguous bytes of a codeword are now spaced further apart in time from each other and interspersed with bytes from a number of other codewords. The process of rearranging the order in which bytes from multiple codewords are transmitted is called interleaving. Because impulse noise is relatively short in duration, spreading out portions of codewords in time ensures that any instance of impulse noise will tend to corrupt a smaller number of bits from any given codeword. See, ITU-T Rec. G.992.3 (01/2005) at p. 38) ("To spread the Reed-Solomon codeword and, therefore, reduce the probability of failure of the FEC in the presence of impulse noise, the FEC Output Data Frames shall be convolutionally interleaved."). The combination of RS coding followed by interleaving is referred to as "interleaving/RS coding." See, e.g., '956 Patent at 18:5.

65. Upon receipt of the bytes of interleaved codewords, the codewords are reassembled by rearranging the bytes back into their original order. The reverse process of reordering/reassembling the codewords to their original order is referred to as deinterleaving. The combination of deinterleaving followed by RS decoding is referred to as "deinterleaving/RS decoding."

66. The illustration below depicts an interleaved data stream consisting of four bytes from four different codewords.[7] In particular, the first byte from the four different codewords (yellow #1a, green #2a, blue #3a, and orange #4a) is transmitted first; then a second byte from

[7] For the sake of simplicity, portions of this discussion of interleaving and deinterleaving use an interleaver block size equal to the size of a full codeword. It should be understood, however, that interleaving/deinterleaving may be performed on interleaver blocks that are a fraction of the size of a full codeword. For example, the interleaver may operate on blocks of bytes that are one-half or one-quarter of a full codeword.

each codeword (yellow #1b, green #2b, blue #3b, and orange #4b) are transmitted and so on. As shown below, impulse noise corrupts only 25% or less of the total number of bits in any codeword. Thus, once the data stream is deinterleaved so that the original codewords are reassembled, the effect of impulse noise is spread across multiple codewords and no individual codeword has more than 25% of its bits corrupted. This would allow the error correction code to recover all of the original data bits.

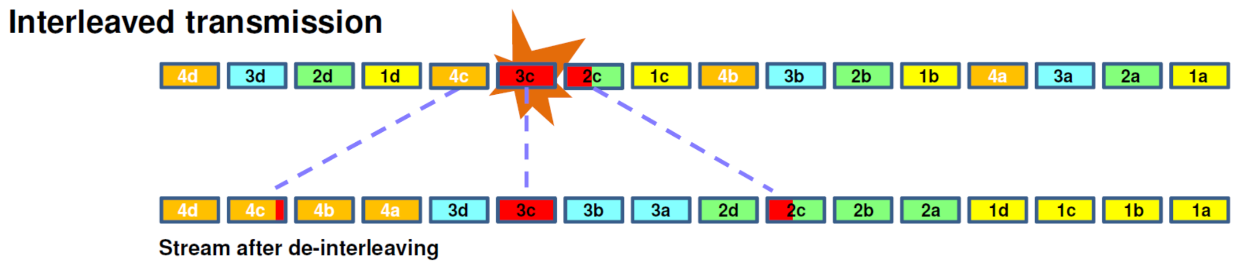

67. Latency is a very important parameter of communications systems. It includes the delay introduced by the protocol used to communicate the data. I will use the terms delay and latency interchangeably.

68. Interleaving increases throughput, because it reduces the need for retransmission. However, it has a downside because, as previously explained interleaving contributes to latency. This is because the all of the symbols/bytes of a codeword must be received before deinterleaving and decoding can be performed. For example, in the illustration above, the first group of yellow bits (yellow #1a) must be stored while waiting for the second group (yellow #1b), third (yellow #1c), and fourth group (yellow #1d) of yellow bits to be received. Only after receipt of all of the bits of a codeword can the codeword be decoded. This delays the recovery of the original data bits. Similarly, on the transmitter side, the last of the codewords within an interleaved block of multiple codewords must be processed (e.g., encoded) before the first group of bits in the interleaved block can be transmitted. The time delay introduced because of the

error correction and coding and interleaving/deinterleaving operation contributes to an increase in the overall latency. This increase in latency is a function of the number of bytes in a codeword (i.e., "codeword size"), the amount by which bytes of a codeword are spread out (i.e., "interleaver depth"), and the data rate.

69. An interleaver is a component in a transceiver that accepts bytes of codewords and outputs the identical bytes but in an order different from the order in which the bytes were received.

70. A deinterleaver is a component in a transceiver that receives a block of interleaved bytes, and reassembles the bytes back into their original order. The deinterleaved codewords may then be decoded to recover the original data bits.

71. A transceiver includes a transmitter and a receiver. Consequently, such a transceiver that uses interleaving and deinterleaving used both an interleaver for transmitting and a deinterleaver for receiving. The interleaver is used to interleave blocks of codewords prior to transmission and the deinterleaver is used to deinterleave received blocks of interleaved codewords.

72. An important element of the interleaver is the interleaver memory that is used to perform the interleaving. As recognized by the Family 3 Patent specification, the size of the interleaver memory needed to perform the interleaving is determined by the interleaver depth and the size of the codewords (or a block of bytes comprising a fraction of a codeword) that are being interleaved. Increasing the interleaver depth increases the immunity to longer bursts of impulse noise, but also increases the amount of memory required to perform the interleaving operation.

73. The simple illustrations above do not actually use the particular type of interleaver

specified by the DSL standards. The DSL standards use a convolutional interleaver. The difference presents itself primarily in the order in which the bytes are rearranged. The operation of a simple convolutional interleaver is illustrated below (with I = 5 and D = 3). "I" is the size of the interleaver block in number of bytes. An interleaver block can comprise a full codeword or a fraction of a codeword. For example, if a codeword is 20 bytes, this could be divided into four interleaver blocks where the interleaver block size "I" would be 5 bytes.

Convolutional Interleaver: Input block length = 5, Depth Parameter D = 3

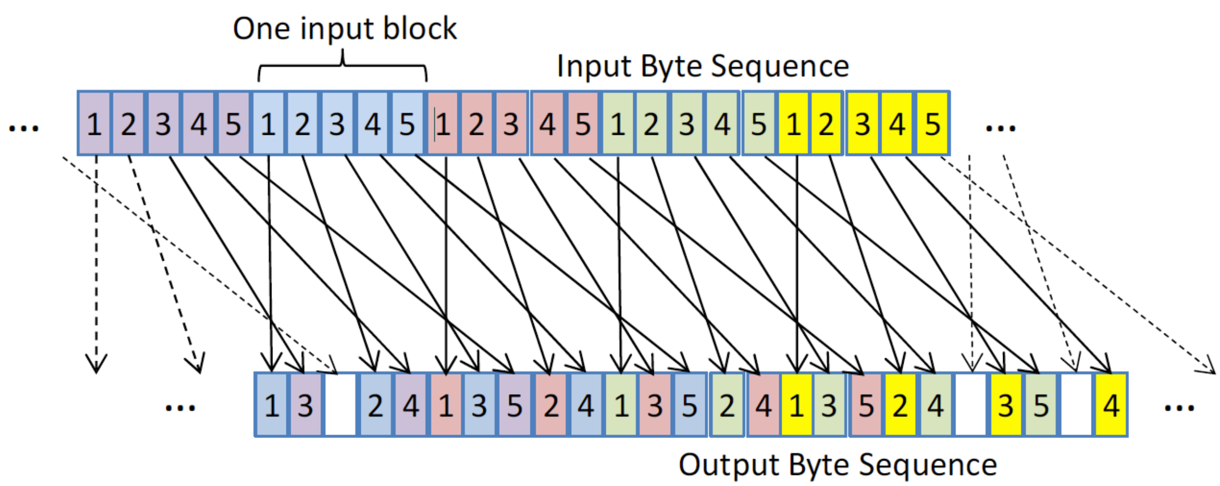

74. A convolutional interleaver works as follows: the first byte[8] of each input block is not delayed in the output; the second byte of each input block is delayed by D-1 bytes; the third byte is delayed by 2*(D-1) bytes; and so on.

75. This results in the last byte of the block (the "I" byte) being delayed (I-1)*(D-1) from its original position. ITU-T Rec. G.993.2 (02/2006) at 51 ("I is the interleaver block size in bytes. Each of the I bytes in an interleaver block B0B1…BI–1 shall be delayed by the interleaver

[8] The term "byte" refers to a grouping of "bits"; for example, it is often used to refer to a group of eight bits. *See*, Hargrave's Communications Dictionary, (2001) at p. 69 (defining "**byte (8)** A contiguous set of bits operated on as a single unit. A byte is usually shorter than a word, and usually 8 bits are grouped as a byte.").

by an amount that varies linearly with the byte index. More precisely byte Bj (with index j) shall be delayed by $\Delta[j] = (D - 1) \times j$ bytes, where D is the interleaver depth in bytes, and D and I are co-prime (have no common divisor except for 1)").

76. The memory requirements for a convolutional interleaver can be determined from two parameters of the interleaver: block length and depth. The illustration above includes a block length of 5 and a depth of 3. The block length will typically be the length in bytes of a codeword, or of a fraction of a codeword (for example, a codeword may be broken up into an integer number of blocks). The depth parameter effectively determines how spread apart adjacent bytes of an input block will be in the output of the interleaver (for example, each of bytes 1 and 2 in the illustration above, which are adjacent in the input block, are spread three bytes apart in the output).

77. Memory is required for interleaving because bytes must be stored in order to rearrange the sequence in which the bytes are output. For example, as illustrated above, the last byte (byte #5) of the blue block is delayed by eight byte locations from where it would have been output but for interleaving and, thus, byte #5 would have to be stored during the delay. The combined delay due to interleaving at the transmitter and deinterleaving at the receiver (i.e., end to-end delay) is a function of the block length "I" and the depth parameter "D" of the interleaver. "The end-to-end delay in octets for the interleaver and de-interleaver pair on path p, with p = 0, 1, is given by: $delay\_octet_{x,p} = (I_{x,p} - 1) \times (D_{x,p} - 1)$ where the direction of transmission x is either 'DS' for the downstream channel or 'US' for the upstream channel, $I_{x,p}$ is the interleaver block length, and $D_{x,p}$ is the interleaver depth." ITU- T Rec. G.993.2 (02/2006) at 23.

78. For one upstream channel and one downstream channel, the aggregate interleaver and deinterleaver delay is the sum of $delay\_octet_{DS,0}$ and $delay\_octet_{US,0}$, which sum is referred to

as MAXDELAYOCTET. The amount of memory required at each transceiver is determined by the value of MAXDELAYOCTET. A694 (ITU-T Rec. G.993.2 (02/2006) at 23) ("The minimum amount of memory required in a transceiver (VTU-O or VTU-R) to meet this requirement is $\frac{\text{MAXDELAYOCTET}}{2}$ octets.").

4. Memory and Memory Sharing

79. As described above, memory is required for the forward error correction (including Reed Solomon coding), interleaving functionality (including deinterleaving), and retransmission. This section will provide an overview of memory and the sharing of memory.

80. Memory is used to store information. IEEE Standard Dictionary of Electrical and Electronics Terms, 4th ed. (1988) at p 582 ("memory. See: storage") and p. 956 (defining storage as "[a]ny device in which information can be *stored*, sometimes called a memory device") (emphasis added); see also Standard Dictionary of Computers and Information Processing, Martin H. Weik, 3rd printing (1970) at p. 186 (stating that memory is same as storage) and p. 271 (defining storage as "a device . . . which receives data, holds and, at a later time, returns data.").

81. Memory is comprised of memory cells. "A memory cell is the smallest subdivision of a memory into which a unit of data has been or can be entered, in which it is or can be stored, or from which it can be retrieved." IEEE 100 The Authoritative Dictionary of IEEE Standard Terms, 7th ed. (2000) at p. 685 (defining memory cell as "the smallest subdivision of a memory into which a unit of data has been or can be entered, in which it is or can be stored, and from which it can be retrieved.") Shown below is an illustration of memory. Each cell of the grid corresponds to an exemplary memory cell.

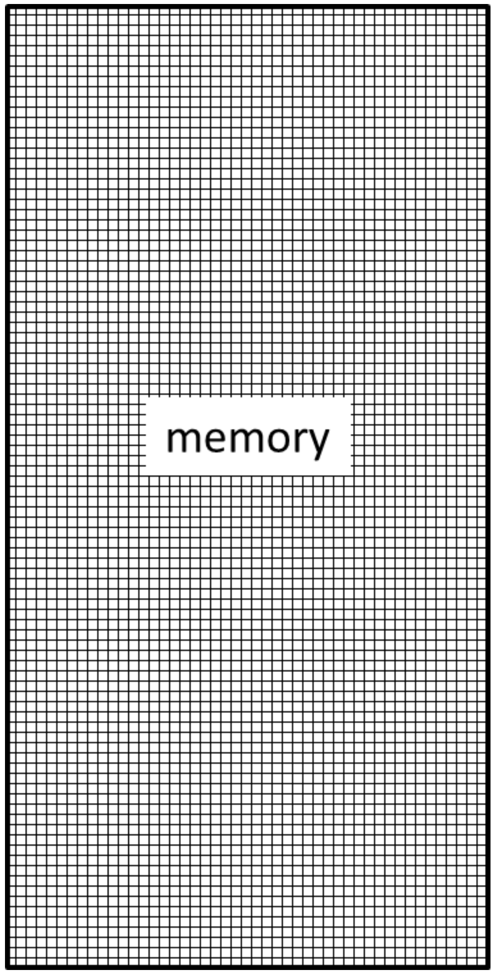

82. The Family 3 Patents recognize that an interleaver and deinterleaver can consume a large amount of memory. *See* '890 Patent at 1:45-48 ("[A]n interleaver within a latency path can consume a large amount of memory in order to provide error correcting capability."). The cost of memory can contribute significantly to the cost of DSL equipment.

83. Practical DSL transceivers are provided with a fixed amount of memory. Inventions described and claimed by the Family 3 Patents relate to "memory sharing in communication systems." '890 Patent at 1:16-17. Specifically, the Family 3 Patents describe schemes to allocate shared memory between an interleaver and a deinterleaver of a transceiver. The interleaver and the deinterleaver of the transceiver each use all or a portion of its allocation of memory to perform interleaving and deinterleaving, respectively.

84. In the context of the Family 3 Patents, shared memory is common memory where portions can be allocated for either interleavering or deinterleaving in a transceiver. *See* ’890 Patent at 1:57-59 (“[M]ore particularly, an exemplary aspect of this invention relates to shared latency path memory in a transceiver.”); ’890 Patent at 5:33-39 (“[F]or example, an exemplary transceiver could comprise a shared interleaver/deinterleaver memory and could be designed to allocate a first portion of the shared memory 120 to an interleaver, such as interleaver 216 in the transmitter portion of the transceiver and allocate a second portion of the shared memory 120 to a deinterleaver, such as 316, in the receiver portion of the transceiver.”)

85. The figure below illustrates an example allocation of the shared memory between an interleaver and a deinterleaver. The portion of shared memory above the dashed line is allocated for use by the interleaver and the portion of shared memory below the dashed line is allocated for use by the deinterleaver. The interleaver uses memory cells from the portion of shared memory allocated to interleaving and the deinterleaver uses memory cells from the portion of shared memory allocated to deinterleaving.

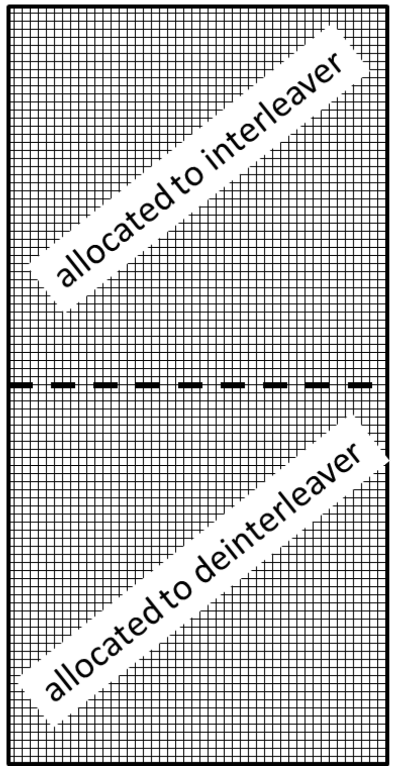

86. The portion of the shared memory allocated to the interleaver can be used at the same time as the portion of the memory allocated to the deinterleaver. For example, the interleaver may be using a portion of the shared memory to store (i.e., hold) bytes of codewords that are being interleaved at the same time that the deinterleaver is using a portion of the shared memory to store (i.e., hold) bytes of codewords that are being deinterleaved. *See* A35 ('882 patent) at claim 13, 12:41-43 ("wherein the shared memory allocated to the deinterleaver is used at the same time as the shared memory allocated to the interleaver.").

87. The sharing of resources such as memory requires coordination between the two transceivers that form one DSL connection. To that end, the Family 3 Patents explain that the allocation may be performed in accordance with messages exchanged by the transceivers. *See, e.g.,* '890 Patent at 8:47-51 ("[I]n this case, the first transceiver would send a message to the second transceiver containing the information described above and based on this information and the application requirements the second transceiver would select latency path settings."). For

example, based on a received message, the transceiver may allocate a first portion of the shared memory for use by an interleaver and/or allocate a second portion of the shared memory for use by a deinterleaver.

88. The Family 3 Patent specification explains that the allocation of shared memory between the interleaver and deinterleaver is not fixed. Instead, "the sharing of resources can be modified and messages transmitted between [] two transceivers at any time during initialization and/or user data transmission." '890 Patent at 4:37-40. In other words, the messages exchanged between the transceivers are used to coordinate initial allocation, and any subsequent changes in the allocation, of shared memory. '890 Patent at 7:54-60 ("This information may need to be forwarded during initialization in order to initially configure the DSL connection. This information may also need to be forwarded during SHOWTIME in order to reconfigure the DSL connection based on a change in applications or the application requirements.").

89. The allocation may be adjusted in response to "a change in communication that would require the adjustment of the shared resource allocation." '890 Patent at 9:29-32. "Examples of changes in communications conditions include a change in applications being transported over the system and/or changes in the channel condition, etc." '890 Patent at 9:32-34. The Family 3 Patent specification explains that applications like "video typically requires a low BER [bit error rate] (<1E-l0) but can tolerate higher latency (>20 ms). Voice, on the other hand, typically requires a low latency (<1 ms) but can tolerate BER (> lE-3)." '890 Patent at 1:27-30. Thus, if the traffic being handled changes from voice to video, the codeword size and the interleaving depth (together with the corresponding deinterleaving depth) may be increased to lower the bit error rate. This results in an increase in the memory size used by interleaving and its corresponding deinterleaving.

90. Within one transceiver, this may require a new allocation of memory so that portions of the shared memory that were previously used by the interleaver are now used by the deinterleaver. Thus, one or more memory cells that had been allocated to the interleaver may be used by the deinterleaver. Such a change in allocation is illustrated below. On the left (Allocation 1), a first allocation is shown with about 35% of the memory cells allocated to the interleaver and about 65% allocated to the deinterleaver. On the right (Allocation 2), a different allocation is shown with about 20% of the memory cells allocated to the interleaver and about 80% allocated to the deinterleaver. Per this example, about 15% of the memory cells that were at one time allocated to the interleaver are at another time allocated to the deinterleaver.

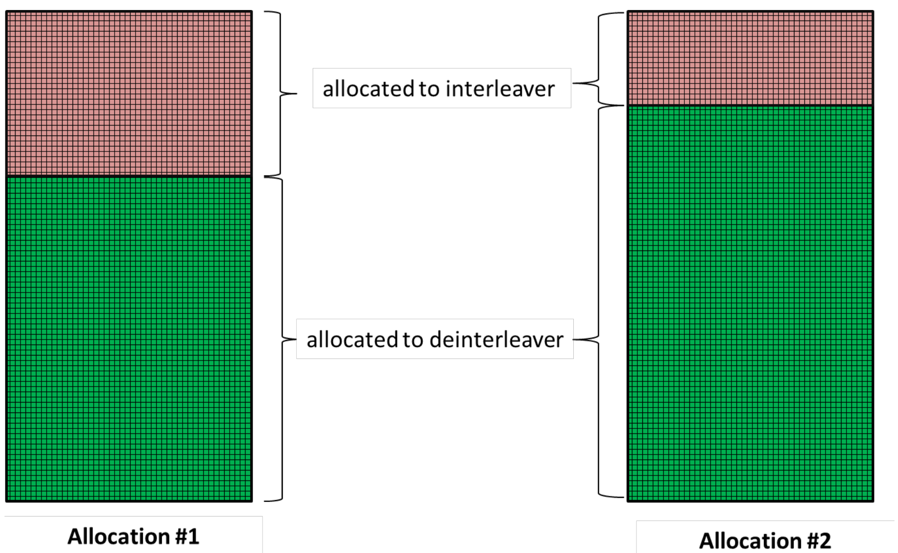

91. Certain inventions claimed by the Family 9 Patents relate to sharing of transceiver memory "between a retransmission function and one or more of the interleaving/deinterleaving functionality, RS coding/decoding functionality." '956 Patent at 22:20-24. The Family 9 Patents explain that "interleaving and RS coding methods [, on one hand,] and retransmission protocols[, on the other hand,] provide different advantages with respect to error correction capabilities, latency, buffering requirements, and the like. For example, under certain configuration and noise

conditions the interleaving/RS coding provides error correction/coding gain with less delay and overhead than the retransmission protocol (for packets that can be retransmitted). While under other conditions the retransmission protocol will provide better error correction with less delay and overhead than the interleaving/RS coding." '956 Patent at 17:63-18:5. Accordingly, the Family 9 Patents contemplate using either interleaving/deinterleaving/RS coding/RS decoding for data communicated on both the upstream and downstream channels, retransmission for both the upstream and downstream channels or a combination of retransmission functionality for either the upstream or downstream channels and interleaving/deinterleaving/RS coding/RS decoding for the downstream or upstream channel.

92. In the context of the Family 9 Patents, shared memory is memory that can be allocated for use by a transceiver's retransmission functionality and the transceiver's interleaving/RS coding and/or deinterleaving/RS decoding. For example, the Family 9 Patents recognize that DSL transceivers may use interleaving and deinterleaving for both the upstream and downstream channels and accordingly require an "aggregate bidirectional interleaver and de-interleaver memory of 65 Kbytes . . . that corresponds to storage requirement of approximately 32 Kbytes in a single transceiver." '956 Patent at 17:45-53. The Family 9 Patents recognize that "under certain configuration and noise conditions the interleaving/RS coding provides error correction/coding gain with less delay and overhead than the retransmission protocol (for packets that can be retransmitted). While under other conditions the retransmission protocol will provide better error correction with less delay and overhead than the interleaving/RS coding."'956 Patent at 17:63-18:5. Accordingly, Family 9 Patents disclose allocating the transceiver memory "between a retransmission function and one or more of the interleaving/deinterleaving functionality, RS coding/decoding functionality." '956 Patent at 22:20-24. "For example, 40% of

the memory could be allocated to the interleaving/deinterleaving/RS coding/RS decoding functionality with the remaining 60% allocated to the retransmission of functionality." *Id.* at 18:26-30.

93. The Family 9 Patents explain that "[a]ssociated with the ability to allocate or partition memory between one or more of the interleaving/deinterleaving/RS coding/RS decoding functionality and retransmission functionality, is the ability to exchange information between transceivers on how to establish this allocation." '956 Patent) at 18:66-19:3. "For example, if the receiving modem contains 100 kBytes of available memory, the transmitting modem could send a message to the receiving modem indicating that 25 kBytes should be allocated to RS coding functionality and 75 kBytes should be allocated to the retransmission functionality. Since the receiving modem generally determines the interleaving/RS coding parameters that are used, the receiving modem could use this information to select parameters, e.g., interleaver depth and codeword size, that would result in an interleaving memory requirement that is no more than the amount indicated in the message." '956 Patent at 19:8-19. "Alternatively, or addition[ally], the receiving modem can send a message to the transmitting modem indicating how much of the available memory is to be allocated to one or more of the interleaving/deinterleaving/RS coding/RS decoding functionality, and how much memory should be allocated to the retransmission functionality." A17 ('956 Patent) at 19:20-25.

94. An exemplary message used to indicate memory requirements upon which the allocation of memory will be based is the O-PMS message. See A848 (ITU-T Rec. G.993.2 (02/2006) at p. 177) (noting that the O-PMS message "also specifies the portion of shared interleaver memory that the VTU-R can use to de-interleave the downstream data stream."). The table below sets forth the format of the O-PMS message, which is transmitted by the VTU-O:

Table 12-46 – Description of message O-PMS

	Field name	Format
1	Message descriptor	Message code
2	MSGLP	1 byte
3	Mapping of bearer channels to latency paths	1 byte
4	B_{x0}	1 byte
5	B_{x1}	1 byte
6	LP_0	Latency Path descriptor
7	LP_1	Latency Path descriptor
8	max_delay_octet$_{DS,0}$ ~~$MaxD_0$~~	3 bytes
9	max_delay_octet$_{DS,1}$ ~~$MaxD_1$~~	3 bytes
10	max_delay_octet$_{US,0}$	3 bytes
11	max_delay_octet$_{US,1}$	3 bytes

ITU-T Rec. G.993.2 (2006)/Amd.1 (04/2007) at p.31.

95. Information in fields 8 through 11 is used to establish the allocation of memory between the interleaving/deinterleaving/RS coding/RS decoding memory in both transceivers that communicate over a communication channel. "Field #8 'max_delay_octet$_{DS,0}$ . . . is a 3-byte field that specifies the maximum interleaver delay that the VTU-R shall be allowed to use to deinterleave the data stream in downstream latency path #0. The maximum interleaver delay shall be specified in bytes as an unsigned integer." ITU-T Rec. G.993.2 (2006)/Amd.1 (04/2007) at p.31; see also ITU-T Rec. G.993.2 (02/2006) at p. 178. "Field #10 'max_delay_octet$_{US,0}$' is a 3-byte field that specifies the maximum interleaver delay that the VTU-O will use to deinterleave the data stream in upstream latency path #0. The maximum interleaver delay shall be specified in bytes as an unsigned integer." ITU-T Rec. G.993.2 (2006)/Amd.1 (04/2007) at p.31.

5. Attenuation, Noise, and Signal-to-Noise Ratio and Margin

96. The following section provides an overview of the field of the invention and the specific subject matter claimed the Family 10 Patents. To facilitate an understanding of the inventions, also included is a description of the state of art.

97. The Family 10 Patents relate to the field of multicarrier modulation. '660 Patent at 1:23–25. More specifically, the Family 10 Patents relates to systems using multicarrier modulation as discrete multitone (DMT) systems. Example DMT systems include, for example, ADSL systems, as specified in the ITU-T Recommendations G. 992.1 (06/99) and G.992.2 (07/99). *See* '660 Patent at 1:45–50. Additional DMT systems are specified in, for instance, the ITU-T Recommendation G.992.3 (07/2002) (hereinafter "G.992.3" or ADSL2) and the very-high-bitrate digital subscriber line (VDSL) systems as specified in the ITU G.993.x standards (e.g., ITUT Recommendation G.993.2 (01/2015).

98. DSL-based systems, as described above, typically operate on twisted pairs of wires. A twisted pair consists of two insulated copper wires of various thicknesses, e.g., 0.016 inch (24 gauge) arranged in a spiral pattern. A major characteristic of the twisted pair is attenuation. Attenuation (i.e., insertion loss), measured in decibels (dB), can be defined as the difference between the transmitted power at one end of the twisted pair and the received power at the other end of the line. Attenuation is mainly determined by wire diameter, line length, and the transmitted signal frequency. The high-frequency components of a transmitted signal will be attenuated more than the low-frequency components.

99. The insertion loss also depends on temperature. The Family 10 Patents recognize that an increase in temperature will affect the higher frequencies more than the lower frequencies. "Therefore, as the temperature increases, in addition to an overall increase in insertion loss, the insertion loss at the higher frequencies increases more than the insertion loss at the lower frequencies." *See* '660 Patent at 4:26–29.

100. A transceiver that communicates data by modulating data bits simultaneously onto multiple carriers is referred to as a multicarrier transceiver. Multicarrier transceivers divide

the frequency bandwidth of the communication link (channel) into multiple carriers (also known as subcarriers, subchannels, or bins). The carriers operate in parallel in the sense that each carrier carries a fraction of the total data stream.

101. For example, the ITU Recommendations for ADSL, such as G.992.1 and G.992.3, separate the entire bandwidth into 255 carriers separated by $\Delta f = 4.3125$ kHz. The modulation process combines the used plurality of carriers into a signal that can be transmitted. In this example, one carrier occupies 4.3125 kHz wide sub-band within the overall bandwidth of the transmitted signal.

102. In operation, the multicarrier transceiver divides the data to be communicated over the communication link into groups of data bits, allocates each group of data bits to a respective carrier, and modulates each group of data bits onto its respective carrier. The size of each group may range typically from zero to 15. That is to say, between zero and 15 bits can be transmitted on each carrier.

103. It is possible to allocate the same number of bits to every carrier. In typical DSL systems, however, different numbers of bits may be allocated to different carriers. The process of determining the number of data bits allocated to a particular carrier is referred to as "bit allocation" or "bit loading."

104. The number of bits allocated to particular carrier depends primarily on the Signal to Noise Ratio (SNR) of that carrier and a maximum error rate requirement (for example, a bit error rate (BER)) for the communication link. *See* '660 Patent at 1:51–54. The SNR on a given subchannel is impacted by how much actual noise is present in the frequency range of the subchannel and the amount of signal attenuation on the subchannel (i.e., the drop in signal power between the transmitter and receiver).

105. SNR is a term that is well understood in the art. The SNR is typically determined by measurements made at the receiver during initialization of the multicarrier transceiver. Measured in dB, the SNR per carrier is generally the difference between the power of the carrier and the power of the actual noise at that carrier measured at the receiver.[9]

106. As discussed above, there are several noise sources including, inter alia, thermal noise, impulse noise, crosstalk from other DSL lines or systems, and the like. *See* '660 Patent at 5:4–7. While there is an amount of noise at all times, the same amount of noise is not present all the time. For example, crosstalk noise may change as adjacent DSL lines are activated or deactivated. The SNR also depends on the attenuation. Since temperature affects attenuation, the SNR may change as temperature changes.

107. A group of data bits are modulated onto a carrier using typically Quadrature Amplitude Modulation (QAM). Furthermore, DSL systems may be required to operate at an error rate that does not exceed a specified maximum rate. For example, certain ADSL systems were required to operate at a BER that did not exceed 1×10^{-7}. It is known in the art that digital modulation techniques need a certain minimum SNR to achieve a required error rate. Bit allocation is generally performed by allocating a higher number of bits to carriers with higher SNR and allocating a fewer number of bits to carriers with lower SNR.

> For example, if the required BER is $1 \mathrm{x} 10^{-7}$, i.e., one bit in ten million is received in error on average, and the SNR of a particular subchannel is 21.5 dB, then that subchannel can modulate 4 bits, since 21.5 dB is the required SNR to transmit 4 QAM bits with a $1 \mathrm{x} 10^{-7}$ BER. Other subchannels can have a different SNR and therefore may have a different number of bits allocated to them at the same BER.

[9] If not measured in dB, the SNR is dimensionless and is generally the ratio (rather than the difference) of these powers.

'660 Patent at 1:55–61.

108. The noise that practical DSL systems experience is time-varying, due to, *inter alia*, changes in crosstalk, impulse noise, and temperature. When there is an increase in the noise level, it is desirable for DSL systems to be able to continue to operate without an increase in error rate. The SNR margin is a configurable parameter that allows this goal to be accomplished. The SNR margin is used to determine the number of bits allocated to each of a plurality of carriers that allow for an increase in noise associated with the plurality of carriers. '660 Patent at 1:65–2:3. The value of the SNR margin specifies an extra SNR requirement per carrier that is in addition to the minimum SNR required to maintain a specified error rate for the communication link. *See* '660 Patent at 1:65–2:3.

109. The ITU Recommendations for ADSL incorporate an SNR margin parameter during bit allocation. Specifically, Section 9.5.1 of G.992.1 explains that the SNR margin represents the amount of increased received noise (in dB) relative to the noise power that the system is designed to tolerate and still meet the target BER of 1×10^{-7} (i.e., one bit in ten million is received in error on average). *See* G.992.1 at p. 85. Multicarrier transceivers use the SNR margin parameter to increase the DSL system's immunity to changing communication link conditions. '660 Patent at 2:11–12, 2:29–33. When a multicarrier transmission system is operating with a positive SNR margin, the noise can increase instantaneously by the level of the SNR margin and the multicarrier transmission system will still maintain the required error rate. *See* '660 Patent at 2:16–19. In other words, a positive SNR margin serves to increases the robustness of the system.

110. A potential penalty for this increase in robustness, however, is a decrease in the data rate. *See* '660 Patent at 2:23–33. The data rate decreases because fewer bits are allocated to

carriers when a positive SNR margin is applied. For example, the required SNR to transmit 4 bits with a $1x10^{-7}$ BER may be 21.5 dB, where the SNR margin is 0 dB. See ’660 Patent at 1:55–59. Increasing the SNR margin from 0 dB to 6 dB would result in a 27.5 dB SNR (i.e., 21.5 dB SNR + 6 dB SNR margin) being required to transmit the same number of bits (i.e., 4 bits). *See* ’660 Patent at 2:19–33. If the DSL system is operating at a 6 dB margin, however, the noise levels can unexpectedly increase by 6 dB and the system will still operate at the required $1x10^{-7}$ BER. *See* A5 (‘660 Patent) at 2:21–33.

111. As can be appreciated, there exists a compromise between robustness (e.g., resilience to noise) and the achievable data rate (e.g., number of bits per carrier). *See* ’660 Patent at 2:27–29. Indeed, a higher SNR margin results in a higher level of immunity to changing channel conditions at the expense of the achievable data rate, whereas a lower SNR margin results in a higher data rate at the expense of a lower immunity to changing channel conditions. *See* ’660 Patent at 2:21–33.

112. Multicarrier transmission systems traditionally settle on a single SNR margin, which is then applied to all carriers for purposes of bit allocation. *See* ‘660 Patent at 2:34–35. The single SNR margin is often selected to achieve an acceptable data rate, while providing an amount of immunity to degraded SNR for each carrier. For example, ADSL systems often use a fixed 6 dB SNR margin on all carriers carrying data bits. *See* ’660 Patent at 2:35–38. Using a single SNR margin for all of the subchannels has certain disadvantages, including that picking a single margin required a compromise in view of the possibly different desired rate and/or robustness goals for different subchannels.

113. The Family 10 Patents provide a novel solution to address the tradeoff between channel robustness and the achievable data rate. The Family 10 Patents describe and claim a

multicarrier transceiver that can assign different SNR margins to different carriers in a multicarrier system. *See* ’660 Patent at 4:11–13. That is, one or more carriers may employ a higher SNR margin to provide increased robustness, while at the same time one or more other carriers may employ a lower SNR margin to achieve a greater data rate. *See* ’660 Patent at 4:13–17. By setting a higher SNR margin for some carriers, a higher level of robustness is achieved for those carriers. Setting a lower SNR margin for other carriers can achieve a higher data rate. *See* ’660 Patent at 2:11–12, 2:29–33.

114. The advantages of the Family 10 Patents include, *inter alia*, higher overall data rate and improved overall robustness. For example, "…[s]ince certain subchannels are not overly burdened with a common margin, the overall data rate of the system can be increased without sacrificing the robustness of the system." *See* ’660 Patent at 5:8 –11.

C. Bonding

115. The inventions in the Family 2 Patents relate to improvements to a technique known as bonding. In the context of DSL, bonding is a technology that uses multiple phone lines to improve the rate of a DSL communications network. There is no electrical connection among the bonded phone lines. In the context of DSL, bonding is a technology that uses two or more phone lines, each of which is used to carry a portion of a single stream of data. The lines are referred to as bonded because each line conveys a separate DSL signal that contains a portion of the total traffic stream. I will refer to the total stream as combined or aggregate stream and will refer to each connection, when part of a bonded connection, as an individual connection, to distinguish it from the aggregate, or, alternatively, combined, connection. To illustrate the basic concept of bonding, I prepared the following Figure:

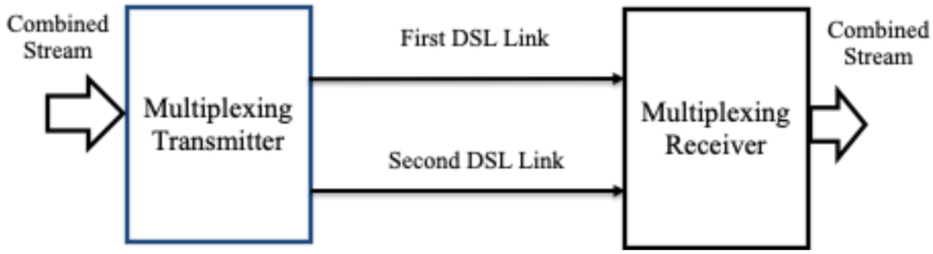

116. As illustrated, in a bonded system, a single stream of data is split by distributing (i.e., de-multiplexing) portions of the data stream on to one or the other of two phone lines and then recombined (i.e., multiplexed) into a single stream once received.

117. The main advantages of bonding are higher data rate and/or reach of a single DSL connection. The data rate of the total stream is approximately the sum of the data rates of the individual DSL connections. For example, if we consider bonding over two DSL links that have data rates of 1 and 1.3 Mb/s, the data rate of the combined stream would be 2.3 Mb/s. The different data rates of the two individual links may be due to the different channel quality conditions (e.g., different SNRs) present on each link.

118. Bonding can be used for different types of traffic, for example such as Asynchronous Transfer Mode (ATM) or Ethernet (also called Packet Transfer Mode (PTM)). ATM is based on asynchronous time-division multiplexing of fixed-length cells, where one cell has a 5-byte header and 48-byte payload. Ethernet uses variable size packets. Bonding for ATM connections is described in ITU-T Recommendation G.998.1 "ATM-based multi-pair bonding." Bonding for Ethernet transport is described in ITU-T Recommendation G.998.2. Bonding may be used to for ADSL or VDSL links.

119. The Family 2 patents describe an example embodiment of bonded transmissions of asynchronous transfer mode (ATM) over ADSL. However, the systems and methods of the Family 2 patents can be practiced using other packet-based protocols and DSL variants.

120. For example, Figure 1 of G.998.1 is reproduced below:

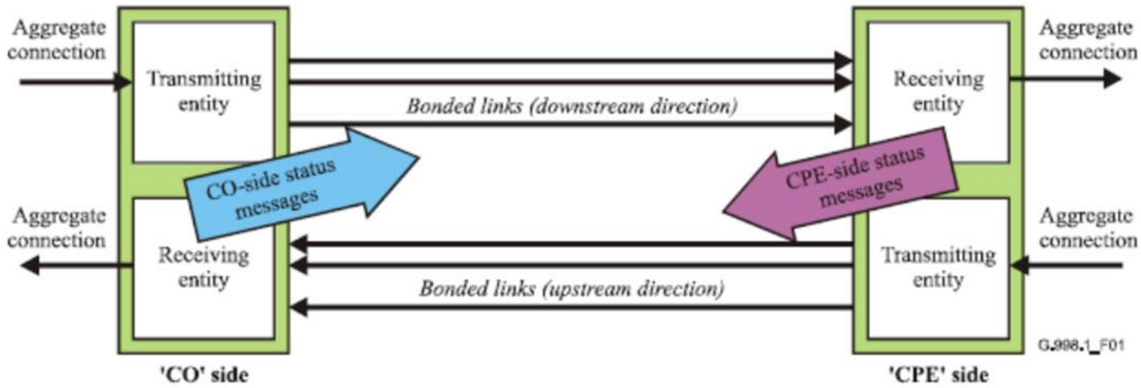

Figure 1/G.998.1 – Status message flow. Each transmission direction is treated independently

ITU-T Recommendation G.998.1 ATM-based multi-pair bonding at page 4.

121. Any communications system consists of three entities: transmitting entity, receiving entity, and channel entity. The transmitting entity and the receiving entity are on different sides of the channel entity. In Figure 1 of G.998.1 these sides are referred to as "CPE side" (Customer Premise Equipment side) and "CO side." In a communications system that uses bonding, the channel consists of multiple links. These links are referred to as "bonded links" in Figure 1 of G.998.1.

122. There are two main special cases regarding the data rates of the two DSL links. The first case is when the data rates of the individual connections are equal. The second case is when the data rates of the individual connections are different.

123. The maximum achievable data rates of the individual links may be different. Implementation would be simpler if the data rates of the individual connections were equal, but this would require setting all data rates equal to the lowest data rate achievable over any of the individual lines and, therefore, result in a lower combined data rate for the bonded line. It is more efficient if the bonded links are allowed to operate with different data rates over the

individual connections.

124. Bonding can also be used to deliver DSL service over a greater distance than possible with a single connection while achieving the same data rate as a shorter non-bonded connection. For example, if the total data rate is kept the same, bonding two lines allows the distance to be extended approximately 1.6 times compared with a single line.

125. Based on the foregoing, bonding requires multiple phone lines, where each phone line is terminated by bonded transceivers (e.g., ATU-Cs or VTU-Os capable of bonding) on the CO side and bonded transceivers (e.g., ATU-Rs or VTU-Rs) on the CPE side.

126. ADSL equipment consists of a transmitting entity and a receiving entity. "[A]s shown in Figure 1, each side of a bonding group has a transmitting and a receiving entity." See ITU-T Recommendation G.998.1 at page 6. The transmitting entity and the receiving entity on the same side of a physical link together comprise a transceiver.

127. Bonding is distinctively different from load balancing. Loading balancing is another technique that also uses multiple phone lines. For example, if there are two computers using two phone lines, the traffic from each computer will transmit and receive over only a single respective phone line. Load balancing does not combine (or aggregate) the individual data streams and does not create a single high-rate stream from multiple lower-rate connections.

128. Figure 2 of the '881 patent illustrates an exemplary bonded system wherein the ATU-C multi-pair multiplexer is marked as 140, the bonded ATU-C transceivers are marked as 150, the bonded lines are marked as 160 and 170, the bonded ATU-R transceivers are marked 210, and the ATU-R mutli-pair multiplexer is marked as 220:

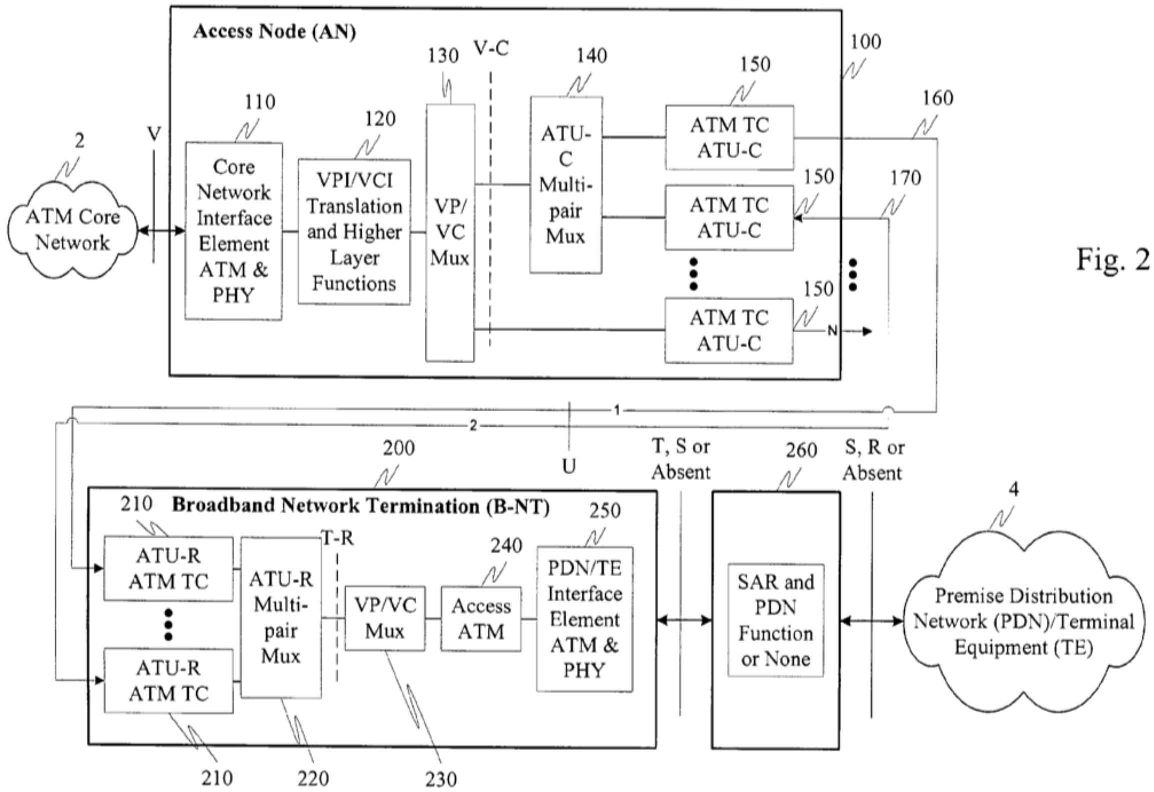

Fig. 2

’881 patent, Fig. 2.

129. Bonding generally involves three distinct processes: (1) Data is split up into separate sub-streams; (2) Those sub-streams are transmitted simultaneously (i.e., in parallel) over the plurality of DSL lines in the bonded connection; (3) The data is recombined into a single stream. In an example embodiment of the ‘881 patent, a “multiplexing transmitter” splits the data into sub-streams and a “multiplexing receiver” recombines the sub-streams into a single data stream. In general, a multiplexing transmitter may also be called inverse multiplexer or de-multiplexer, and a multiplexing receiver may be called simply a multiplexer.

130. A multiplexing transmitter performs a function that can be described as “one-to-many”, i.e. from one stream multiple sub-streams are produced. This process may be thought of as distributing the single stream. This process is illustrated in Figure 3 of the ’881 patent:

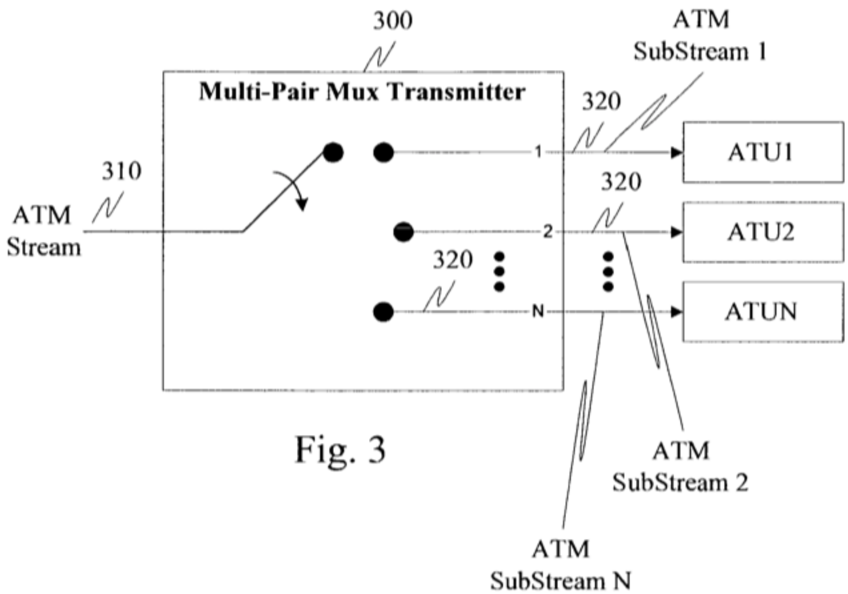

Fig. 3

’881 Patent, Fig. 3 and col. 4:57-58 (“FIG. 3 illustrates an exemplary multi-pair multiplexing transmitter according to this invention”).

131. After the individual sub-streams are produced by the multiplexing transmitter, they are sent over the individual ADSL connections. Of particular importance for the Family 2 patents is that the individual ADSL connections will deliver the sub-streams with different delays.

132. On the other end of the communications system, the data distributed on to the sub-streams must be recombined. A multiplexing receiver performs a function that can be described as “many-to-one”, i.e., from the multiple sub-streams it produces one combined stream. This process is illustrated in Figure 6 of the ’881 Patent:

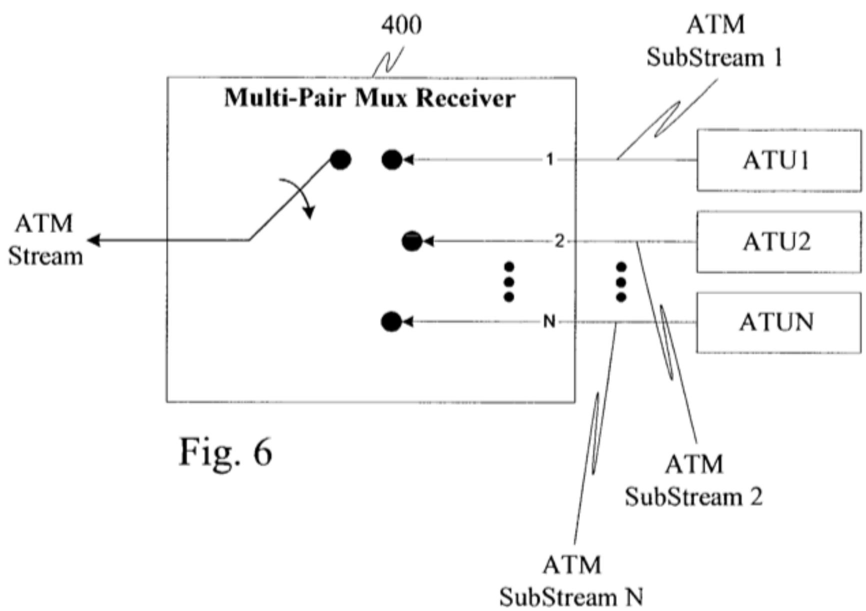

Fig. 6

’881 Patent, Fig. 6.

133. A key consideration is that the multiplexing receiver must be able to recreate the original aggregate stream of data even if the individual data bits are not received in the same order as they were transmitted in the original stream of data.

134. As described above, latency (or, which is the same, delay) is a very important parameter in communication systems. There are several reasons why the different DSL links will achieve different delays. These reasons include, for example, different data rates and different coding parameters such as codeword size and interleaving. The ’881 Patent recognizes that the individual connections that are part of a bonded group might have different latencies: “[F]or example, some of the exemplary reasons for having different delays over different DSL PHYs include, but are not limited to, configuration latency which is based on the configuration of the DSL transmission parameters. Specifically, these parameters include the data rate, coding parameters, such as the coding method, codeword size, interleaving parameters, framing parameters, or the like.” ’881 Patent, col. 6:10-16.

135. The ’881 Patent also identifies the problems that might arise due to the difference

in latencies among the individual bonded links: "[T]his potential latency difference between bonded PHYs places implementation requirements on the multi-pair multiplexer. In particular, the multipair multiplexer receiver must be able to reconstruct the ATM stream even if the ATM cells are not being received in the same order as they where [sic] transmitted." '881 Patent, col. 6:4-9. In other words, if cells (or packets) received via one bonded link are subject to greater latency than the cells received from another bonded link, the original order of transmission must be reconstructed.

136. One way to compensate for a difference in latency between bonded transceivers is for the multiplexing receiver to use a large buffer to store cells or packets from the bonded link with lower latency while awaiting receipt of the cells or packets from the bonded link with higher latency. This is disadvantageous, however, at least because a larger buffer increases the size of the memory and, therefore, the cost of DSL equipment. Other problems associated with high differential latencies among bonded links, including that the mechanism for keeping track of cell counts for reordering becomes more complex.

137. The '881 Patent teaches that an "effective method of reducing the difference in latency between DSL PHYs is mandate that all DSL PHYs are configured with transmission parameters in order to provide the same configuration latency. An exemplary method of accomplishing the same configuration latency is by configuring the exact same data rate, coding parameters, interleaving parameters, etc. on all DSL PHY s. Alternatively, different PHY s can have, for example, different data rates but use the appropriate coding or interleaving parameters to have the same latency on all the bonded PHYs." '881 Patent at 6:56-65.

138. In a bonded connection, each of the multiple ADSL links can use different coding and interleaving parameters. The '881 patent recites an example for using different coding and

interleaving parameters as follows.

> As an example, for Reed Solomon coding and interleaving functions as defined in ADSL standards G.992.1 and G.992.3, incorporated herein by reference in their entirety, the latency due to these functions is defined as:
>
> Latency=$N*D/R$,
>
> where N is the number of bits in a codeword, D is the interleaver depth in codewords and R is the data in bits per second.
>
> For example if N=1600 bits, i.e., 200 bytes, D=64 codewords and R=6400000 bps then:
>
> Latency=1600*64/640000=0.016 seconds.
>
> Therefore if, for example, two PHYs have different data rates, R1 and R2 then, in order to bond these PHYs together and have the same configuration latency set:
>
> $N1*D1/R1=N2*D2/R2$,
>
> where N1 and N2 are the bits in a codeword for each PHY and D1 and D2 are the interleaver depths for each PHY.
>
> This can also be rewritten as:
>
> $N1*D1=(R1/R2)*N2*D2$.
>
> Thus, in general, the N1, N1, N2 and D2 parameters must be chosen to satisfy the above equations and this can be accomplished in several ways.
>
> For example, if the configuration latency is specified as 0.016 seconds, and R1=6400000 bps and R2=1600000 then, as described in the example above, N1 anss D1 can be configured as N1=1600 and D1=64. Therefore:
>
> $N2*D2=(R2/R1)*D1*N1=(1600000/6400000)*1600*64=1600*64/4$.
>
> Therefore, for example, N2 and D2 can be configured as (N2=1600, D2=16) or (N2=400, D2=64) or (N2=800, D2=32), etc.
>
> Obviously the same methods can be applied to more than 2 PHYs with different data rates.

’881 Patent, col. 6:64-7:36.

VI. OPINIONS

139. It is my understanding that Defendants allege that a number of terms are

indefinite.[10] I address each of these terms below.

140. I have been instructed that a claim term is indefinite if it fails to "inform, with reasonable certainty, those skilled in the art about the scope of the invention." *BASF Corp. v. Johnson Matthey Inc.*, 875 F.3d 1360, 1365 (Fed. Cir. 2017) (quoting *Nautilus, Inc. v. Biosig Inst., Inc.*, 572 U.S. 898, 901 (2014)).

A. "Reduce a Difference in Latency" Terms (Family 2)

141. The "reduce a difference in latency" terms, to one of ordinary skill in the art, would have been understood with reasonable certainty.[11]

142. The Family 2 Patents show that those of skill in the art would have understood this term to mean refer to a "reduce a difference in configuration latency."

143. This construction clarifies that the "latency" of this claim term refers to ***configuration*** latency. While one of skill in the art would understand that total latency may be attributed to other factors in addition to configuration latency, the claims refer to latency that may only be controlled through transmission parameter values. E.g., '881 patent at claim 1 ("utilizing at least one transmission parameter value, for each transceiver in a plurality of bonded transceivers, to reduce a difference in latency between the bonded transceivers"). The specification also confirms that configuration latency and, thus, a difference in configuration

[10] These terms include "reduce a difference in latency between the bonded transceivers," "each bonded transceiver utilizing at least one transmission parameter value to reduce a difference in latency between bonded transceivers," "utilize at least one transmission parameter value, for each transceiver in a plurality of bonded transceivers, to reduce a difference in latency between bonded transceivers," and "utilize at least one parameter associated with operation of at least one of the first and second transceivers to reduce a difference in latency between the first and second transceivers" ('881 Patent, Claims 17, 25, 26, 29, 31, 33, 37; '193 Patent, Claim 13; '601 Patent, Claim 14, 21.

[11] These "PTM-TC codeword" terms are found in the '577 Patent, Claim 16, 37; '348 Patent, Claims 1, 9; '055 Patent, Claim 17.

latency between bonded transceivers, is controlled by transmission parameter values. ’881 Patent, col. 6:10-16. (“[F]or example, some of the exemplary reasons for having different delays over different DSL PHYs include, but are not limited to, configuration latency which is based on the configuration of the DSL transmission parameters. Specifically, these parameters include the data rate, coding parameters, such as the coding method, codeword size, interleaving parameters, framing parameters, or the like.”).

144. The “reduce a difference” portion of the claim is also understood with reasonable certainty to refer to a reduction in the difference in latency as compared to what the difference could have been but for the utilization of at least one transmission parameter value to reduce the difference. For example, in the example embodiment described in the ’881 Patent at column 6, line 64 through column 7, line 36, the two bonded transceivers each have a different data rate but the transmission parameter values N (codeword size) and D (interleaver depth) for each transceiver is configured so that the difference in latency is reduced.

B. The “Using” Terms (Family 9)

145. The “using” terms, to one of ordinary skill in the art, would have been understood with reasonable certainty.[12]

146. The “using” terms, to one of ordinary skill in the art, would have been understood with reasonable certainty. These terms quite simply specify that functionality of a transceiver (transmitting, for example) is performed using a specified sub-function or sub-element of the

[12] This term is “[transmit / transmitting / retransmit / retransmitting / receive / receiving] [by the transceiver] [at least one packet / a packet / the packet / a retransmitted packet / a message / a plurality of messages / at least one message] using [interleaving / deinterleaving / (a/the) forward error correction encoder / (a/the) forward correction decoder / forward correction encoding / forward correction decoding] [and (an/the) interleaver / and (a/the) deinterleaver / and interleaving / and deinterleaving].” It is included in ’577 Patent, Claims 16, 30, 37, 38, 53, 54; ’348 Patent, Claims 1, 3, 9, 11; ’4473, Claims 1, 3; ’809 Patent, Claims 1, 3, 8, 10, 15, 17, 22, 24.

transceiver (interleaving, for example). This is a natural and common way for one of skill in the art to refer more generically to a function/element and then more specifically to sub-functions/sub-elements that form part of the function/element. Thus, for example, claim 16 of the ’577 patent, which recites “receive at least one packet using deinterleaving,” will be infringed when the subject multicarrier transceiver is operable to perform the function of receiving at least one packet where deinterleaving is performed as part of the receiving. On the other hand, claim 16 would not be infringed if the multicarrier transceiver is not operable to perform deinterleaving as part of the receiving.

C. “Receive a First Plurality of Bits on the First Plurality of Carriers Using a First SNR Margin; Receive a Second Plurality of Bits on the Second Plurality of Carriers Using a Second SNR Margin” (Family 10)

147. The “receive a first plurality of bits on the first plurality of carrier using a first SNR margin” and “receive a second plurality of bits on the second plurality of carriers using a second SNR margin” terms, to one of ordinary skill in the art, would have been understood with reasonable certainty.[13]

148. As I explained above, SNR margin refers to a parameter that is used in the process of bit loading. The SNR margin is used to determine the number of bits allocated to each of a plurality of carriers that allow for an increase in noise associated with the plurality of carriers. ’660 Patent at 1:65–2:3. The value of the SNR margin specifies an extra SNR requirement per carrier that is in addition to the minimum SNR required to maintain a specified error rate for the communication link. *See* ’660 Patent at 1:65–2:3. For example, if two carriers have the same measured SNR and the same SNR margin is used for bit loading for both carriers, generally the same number of bits will be allocated to those carriers. If two carriers have the same measured

[13] This term is found in ’354 Patent, Claim 10.

SNR and a first SNR margin is used for bit loading for the first carrier but a second SNR margin is used for bit loading for the second carrier, generally a different number of bits will be allocated to the first carrier than are allocated to the second carrier.

149. The term "receive a first plurality of bits on the first plurality of carrier using a first SNR margin" is reasonably understood to mean that the claimed multicarrier transceiver is operable to receive a first plurality of bits on a first plurality of carriers where a first SNR margin is used for performing bit loading on the first plurality of carriers. The term "receive a second plurality of bits on the second plurality of carriers using a second SNR margin" is reasonably understood to mean that the claimed multicarrier transceiver is operable to receive a second plurality of bits on a second plurality of carriers where a second SNR margin is used for performing bit loading on the second plurality of carriers.

D. "A Multicarrier Communications Transceiver Operable to Receive a Multicarrier Symbol Comprising a First Plurality of Carriers" (Family 10)

150. One of ordinary skill in the art would have understood what this term means.[14]

151. In general, multicarrier modulation divides the transmission frequency band into multiple subchannels, known as carriers or bins, with each carrier individually modulating a bit or a collection of bits. A transmitter modulates an input data stream containing information bits with one or more carriers, i.e., bins or subchannels, and transmits the modulated information. A receiver demodulates all the carriers in order to recover the transmitted information bits as an output data stream. One of ordinary skill would have understood that a multicarrier symbol that comprises a first plurality of carriers to refer to that multicarrier signal, viewed as the collection of bit steams from a digital perspective (*e.g.*, the collection of bits on each carrier).

[14] This term is found in '354 Patent, Claim 10.

E. "Wherein the First SNR Margin Provides More Robust Reception than the Second SNR Margin" (Family 10)

152. One of ordinary skill in the art would have understood what this term means.[15] One of ordinary skill would have understood that the first signal-to-noise ("SNR") margin provides more robust reception than the second SNR margin by, for example, reducing the bit error rate (*e.g.*, making it less likely that a transmission would be subject to significant errors, which makes the reception more "robust," *e.g.*, less likely to require error correction or retransmission).

I declare under penalty of perjury that the above is true and correct, to the best of my knowledge. Executed on this 15th day of March 2022, in Fort Wayne, Ind.

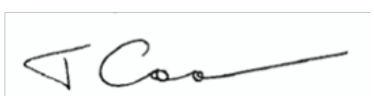

Todor Cooklev, Ph.D.

[15] This term is found in '354 Patent, Claim 10.